

# Sparseness structure of the Gaia least-square problem and the (non-)feasibility of a direct method

Alex Bombrun

Astronomisches Rechen-Institut (ARI)  
European Leadership in Space Astrometry (ELSA)

GAIA-C3-TN-ARI-APB-001-01

February 8, 2008

## 1 Introduction

The Gaia mission has been designed to measure the positions, the velocities and the parallaxes of around one billion sources, mainly stars belonging to our galaxy. Gaia will be an hightech spinning spacecraft, equipped with two telescopes, that share the same focal plane composed of a large CCD panel. Each star will pass many times through the focal plane. In order to build the celestial map with the required accuracy, a few micro-arcsecond, it has been demonstrated that we have to recover the attitude of the spacecraft and calibrate the instruments with the same accuracy. It can only be obtained considering the strong correlation between all the unknowns through all the observations. Gaia will be calibrated by adjusting all the parameters to fit all the measurements as best as possible.

From a mathematical point of view, this problem has been formulated as a least-square system. In this note we investigate its sparseness structure. It appears that the sparseness structure is directly connected with the scanning law and the design of Gaia. In order to estimate the feasibility of a direct computation to solve this least-square problem, we consider a simplified problem and claim that an efficient direct algorithm should first work for the simplified problem. Unfortunately, we will see that a classical direct method for sparse least-square problems, is not efficient on the simplified problem.

## 2 The observation and normal equation for Gaia

### 2.1 Observations, sources and attitude

After the raw-data processing, the main observations made by Gaia can be considered as the star CCD transit times. Gaia measures each time a star crosses a CCD on the focal plane. Each CCD time transit is associated to

a star, some attitude, calibration and global parameters. In this note, since the global and calibration parameters are less numerous than the source and attitude parameters, and in order to simplify the description of the problem, we will assume that the global parameters are perfectly known and that the instruments are perfectly calibrated. Within this assumption, only the source and attitude parameters remain to be determined. The source parameters will be denoted by  $s$  in the following and the attitude parameters by  $a$ . Whereas the source parameters,  $s$ , are independent of the time, the attitude parameters,  $a$ , are only valid on some time intervals. This is an important point to understand the sparseness structure of the least-square problem.

In order to emphasize that the attitude parameters are valid only on a time interval, we introduce a new notation. Let  $\tau$  be a time interval,  $a^\tau$  are all the attitude parameters valid during the time interval  $\tau$ . Similarly,  $t_i^\tau$  denotes all the transit time  $t_i$  associated to the star  $i$  belonging to the interval  $\tau$ . Since we will only consider the correlation between source and attitude parameters, the non-linear observation equation is, when we neglect the noise, the over-determined system of non-linear equations

$$t_i^\tau = f(s_i, a^\tau). \quad (1)$$

where  $f$  is a smooth non-linear function that describes the relation between the observation  $t_i^\tau$  and the unknowns  $s_i$  and  $a^\tau$ .

To give an idea of the number of unknowns and discuss their correlation, we give rough approximations. Actually, the distribution of stars is non-isotropic. We assume that there will be around  $10^8$  primary sources observed by Gaia over five years. Primary sources are sources that have been selected for the good precision of their measurement and their astronomically “benign” nature. In particular, they should have been observed a sufficient number of times. Noting that the precession and spin rate of Gaia have been designed such that each point of the sky is observed 6 times a year and that the overlap of successive scans implies that each primary sources is observed at least 2 times consecutively by the same telescope, each star will transit in average 80 times through the focal plane. That means 720 CCD transit times, observations, per star.

The number of transits influences directly the structure of the observation equation. Indeed the observation equation has a line per CCD time transit. Within our assumption of one hundred million stars, there will be  $7.2 \cdot 10^{10}$  transit times over the whole mission. It represents an average rate of, approximately, 400 transit times per second. That means that around 45 stars are observed each second, since there are 9 consecutive transit times for each star due to the structure of the CCDs panel.

In order to obtain the accuracy expected for Gaia, at least 3 scalar attitude parameters should be considered for each ten seconds of observation. That is approximately  $2.6 \cdot 10^4$  attitude parameters per day,  $5 \cdot 10^7$  over five years. At this point we should note that a B-spline decomposition has been chosen in order to implement an efficient iterative updating of the attitude unknowns, i.e the Astrometric Global Iterative Solution. Indeed as emphasized in [8], splines “are *local* in the sense that a change in the data at point  $t'$  has little effect on the fitted spline at  $t$ ”. This choice has a direct influence on each line of the observation equation. It determines the number of attitude parameters that contribute to each equation (1). Assume that there is a spline knot at each time

$n \in \mathbb{N}$  and that the quaternion  $q^n$  is associated to the spline knot  $n$ , then the equation (1) becomes

$$t_i^n = f(s_i, q^{n-1}, q^n, q^{n+1}, q^{n+2}). \quad (2)$$

It seems reasonable to ask if splines are still adapted if we want to compute a direct solution for the astrometric data reduction. We should be aware that Floor van Leeuwen used a dynamical model of the spacecraft in order to compute a new data reduction for Hipparcos. A discussion with Ulrich Bastian convinced me that the attitude function, which we would like to approximate, is essentially random over timescales of more than a minute or so. We want to estimate the attitude of the spacecraft with an accuracy of a few micro-arcsecond. To obtain this accuracy, the dynamical physical model should probably include the non rigidity of the spacecraft, a measure of the solar pressure, of the solar magnetic field, of the electric currents inboard... certainly much more parameters than the number of observations and the number of sampling spline nodes.

## 2.2 Least square problem

The system (1) is over-determined: there are more equations than unknowns. Due to measurement errors, there does not exist a solution that solves simultaneously all these equations. However the problem becomes mathematically well posed when we try to minimize the norm of the residual vector  $\sum_{i,\tau} \|t_i^\tau - f(s_i, a^\tau)\|_2$ . This is classically a least-square problem.

Since the function  $f$  is non-linear but smooth, we will use a linear approximation around a nominal value  $(\overline{s_i}, \overline{a^\tau})_{i,\tau}$ . Let  $h = (h_{s_i}, h_a^\tau)_{i,\tau}$  be the displacement around that value. We get an over-determined system of linear equations in  $h$ , namely the observation equation

$$l_i^\tau = S_{i,\tau} h_{s_i} + A_{i,\tau} h_a^\tau, \quad (3)$$

with

$$l_i^\tau = t_i^\tau - f(\overline{s_i}, \overline{a^\tau}), \quad (4)$$

$$S_{i,\tau} = \frac{\partial f}{\partial s}(\overline{s_i}, \overline{a^\tau}), \quad (5)$$

$$A_{i,\tau} = \frac{\partial f}{\partial a}(\overline{s_i}, \overline{a^\tau}), \quad (6)$$

Let  $x$  be the vector of unknowns,  $b$  of observations. The observation equation (3) can be written using matrix notation

$$O x = b.$$

The least-square problem is to find a vector  $x$  that minimizes the error  $\|O x - b\|_2$ . A classical way is to form the normal equation,  $O^T O x = O^T b$ . Then, if the normal matrix  $O^T O$  is invertible, the solution to the least-square problem is simply

$$x = (O^T O)^{-1} O^T b. \quad (7)$$

The astrometric data reduction consists in computing such a solution for a selection of primary sources: single stars that have been observed a sufficient

number of times and are astronomically “well-behaved”. Note that this process should be iterate in order to overcome the non-linearities in the equation (1). In the following we will only consider the linear equation (3). There are two kinds of methods to solve a linear equation: direct methods that achieve the computation in a finite number of steps, and iterative methods that are designed to go closer to the solution at each steps. The number of unknowns (see Subsection 2.1) makes this task difficult considering the computer capacity nowadays. The brute-force computation of the inversion of a square matrix with  $10^9$  lines requires approximately  $10^{27}$  basic operations. On the most powerful actual supercomputer, running at 450 teraflops, it will take more than 70,000 years. Hopefully, the Gaia least-square problem being very sparse, it seems better to look for an iterative scheme. Indeed, there are efficient iterative schemes that converge to the solution with a relatively fast rate, see the Hipparcos data reduction [5] and the Astrometric Global Iterative Solution [1]. In the following, we will study the feasibility of a direct method to compute the matrix product (7) considering its sparseness structure.

### 3 Sparseness structure

In this section we will investigate the sparseness structure of the observation matrix  $O$ . The sparseness structure of the observation equation is due to the design of Gaia: two telescopes separated by a large fixed angle and both orthogonal to a spin axis. The scanning law has been optimized to maximize the sky coverage. It leads to a rigid body dynamic with a precession rate of approximately 4 degrees per day and a spin rate of 1 degree per minute [6].

#### 3.1 The block angular structure

From the equations (3) we see that the source equations are independent from each other. Sorting all the observation equations (3) by sources and collecting them in one matrix we get what is called a non-square block angular matrix [3] or, as pointed out in [7], Helmert blocking <sup>1</sup> :

$$\left[ \begin{array}{ccc|c} S_1 & 0 & 0 & A_1 \\ 0 & \ddots & 0 & \vdots \\ 0 & 0 & S_n & A_n \end{array} \right] \begin{pmatrix} h_{s_1} \\ \vdots \\ \frac{h_{s_n}}{h_a} \end{pmatrix} = \begin{pmatrix} l_1 \\ \vdots \\ l_n \end{pmatrix} \quad (8)$$

with  $n$  the number of primary sources ( $10^8$ ), the matrices and vector  $S_i$ ,  $A_i$ ,  $l_i$  the concatenation of, respectively, all  $S_{i,\tau}$ ,  $A_{i,\tau}$ ,  $l_i^\tau$  and, the vector  $h_a$  of all attitude displacement unknowns  $h_a^\tau$ .

Let

$$O = \left[ \begin{array}{ccc|c} S_1 & 0 & 0 & A_1 \\ 0 & \ddots & 0 & \vdots \\ 0 & 0 & S_n & A_n \end{array} \right].$$

we want to compute the solution of the least-square problem  $O h = l$ . We assume that the partial normal matrices,  $S_i^T S_i$ , are invertible (this is guaranteed by the

<sup>1</sup>“Named after the German geodesist F.R. Helmert, who in 1880 published a book in which the method is described.”

choice of primary stars). The normal equation  $O^T O h = O^T l$  is equivalent to the system of linear equations formed by the  $n$  equations

$$\forall i \in [1, 2, \dots, n], \quad S_i^T S_i h_{s_i} + S_i^T A_i h_a = S_i^T l_i,$$

and

$$\sum_{i=1}^n A_i^T S_i h_{s_i} + \sum_{i=1}^n A_i^T A_i h_a = \sum_{i=1}^n A_i^T l_i.$$

A strict forward computation<sup>2</sup> shows that the solution of the normal equation  $O^T O h = O^T l$  can be computed by first solving the reduced normal equation for the attitude displacement,

$$\left[ \sum_{i=1}^n A_i^T S_i (S_i^T S_i)^{-1} S_i^T A_i \right] h_a = \left( \sum_{i=1}^n A_i^T S_i (S_i^T S_i)^{-1} S_i^T l_i \right), \quad (9)$$

and then, forwarding the solution  $h_a$  to solve all the sources equations

$$h_{s_i} = (S_i^T S_i)^{-1} (S_i^T l_i + S_i^T A_i h_a). \quad (10)$$

This structure has been noticed by Hans Bernstein and is used to achieve the one-day astrometric solution [2]. It has been used for Hipparcos in order to collect all the rings and compute the sphere solution [5].

### 3.2 Sparseness of the reduced normal matrix

In this section we analyze the structure of the reduced normal matrix

$$\sum_{i=1}^n A_i^T S_i (S_i^T S_i)^{-1} S_i^T A_i, \quad (11)$$

computed by the process described in Section 3.1 from a list of simulated Gaia-1 observations over 100 days provided by Stefan Jordan. Note that Gaia-1 was designed with a spinning rate of 8 rotations a day, whereas, in the final design, Gaia has a spinning rate of 4 rotations a day. Instead of computing the full matrix product (11), we can more easily estimate the resulting structure. Remember that attitude unknowns are associated to a sampling of the time. We divided each day in 1000 intervals, and, for each couple of intervals,  $(t_1, t_2)$ , we collected the number of stars that have been observed in both intervals. Hence, a zero means that no star has been observed both during the interval  $t_1$  and during the interval  $t_2$ . If, on the contrary, one or more stars have been observed both in  $t_1$  and in  $t_2$ , the corresponding matrix element will be non-zero, represented by dot in Figure 1. Of course, each dot should be considered as a square matrix with the size equal to the number of attitude unknowns per time intervals.

On Figure 1, we plot the sparseness structure of the day 1 (the non-zero values are in black and the zero in white). This figure is similar to Figure 3 in [2]. We can observe the effects of the Gaia design. The two fields of view and the spinning scanning law introduce a banded periodic pattern. The first and

<sup>2</sup>Note that  $S_i (S_i^T S_i)^{-1} S_i^T \neq I$  because  $S_i$  is not a square matrix.




Figure 1: Reduced normal matrix: rows day 1 and columns day 1.

the second band parallel to the diagonal are due to the second field of view, whereas the third band is due to the spinning rotation. The distance between the diagonal and the third band correspond to a full rotation of Gaia around its spin axis, whereas the distance between the first band and the diagonal correspond to the angle between the two fields of view. Moreover, we observe that the farther (from the diagonal) a band is, the sparser the band becomes. This is due to the fact that in addition to the spinning motion, Gaia is rotating around its precession axis. Some points on the celestial sphere are scanned during more than a day whereas some are only visible during a few hours.

On Figure 2, we plot the sparseness structure of the reduced normal matrix between the day 1 and the day 10. We observe a periodic pattern composed of eight points. In order to understand this pattern, we should look at it as two squares, one following the other in time, i.e. in the diagonal. On the scale of one day, Gaia will mainly scan the sky, i.e. the celestial sphere, along a great circle. Two great circles intersect each other in two points symmetric according to the center of the celestial sphere. To each of these points corresponds a square of non-zero values in the reduced normal matrix. The points are transformed in squares because Gaia has two fields of view. The pattern appears periodically because the spin rotation is stabilized at a nominal value. Indeed, the ratio between the edge of each square and the periodicity can give you an approximation of the angle between the two fields of view.

This property is satisfied for all the day-by-day correlations. But, unfortunately for our purpose here, i.e. the computational complexity, Gaia is not only spinning, but also slowly precessing around the Sun-Earth axis, which itself is rotating. Hence, this pattern is no longer periodic on larger timescales than one




Figure 2: Reduced normal matrix, rows day 1 and columns day 10, the filled square isolates a pattern.

day, but slowly moving, and, the points of the pattern are made of 1, 2 or 3 dots. On Figure 3, we superpose the sparseness structure of the first day with day 10 and day 90. Of course the path followed by the pattern is not random but is associated to the scanning law. In an ideal case, if Gaia is only precessing step by step once a day, that means that each day Gaia is scanning along one great circle of the celestial sphere, we should be able to reconstruct the precession law and to predict the location of all the non-zeros of the reduced normal matrix knowing the correlation between two days with all the other. Moreover, time to time, the spinning axis is pointing again toward one of its previous direction. On Figure 4 and Figure 5, we can observe the consequence of Gaia precession law on the reduced normal matrix. Out of the diagonal, the pattern can become similar to the structure observed along the diagonal. The correlation does not occur during more than a few hours because the precession axis are different. Note that the picture on Figure 4 is well oriented. In that case, the strip are not parallel to the diagonal.

### 3.3 Sparsity rate of the reduced normal matrix

We observe that, out of the diagonal, the reduced normal matrix can be decomposed in one pattern of two square that appears *almost* periodically. Let  $\alpha$  be the number of non-zero elements in the pattern, let  $n_r$  the number of spin rotation per day and  $n_d$  the number of day. Then the number,  $N$ , of non-redundant




Figure 3: Reduced normal matrix: rows day 1 and columns days 10 and 90.



Figure 4: Reduced normal matrix: rows first spin rotation of day 1 and columns first spin rotation of days 108.



Figure 5: Reduced normal matrix: rows first spin rotation of day 48 and columns first spin rotation of days 163.

non-zero elements in the reduced normal matrix is of the order

$$N \sim \frac{1}{2} \alpha n_r^2 n_d^2. \quad (12)$$

The number  $\alpha$  of non-zero term in each pattern is directly connected to the equation (2). In the worst case, each corner of the pattern is composed of 3 dots, but some may be empty. Each dot should be consider as a square matrix. The equation (2) results of the choice to approximate the unknown attitude using B-splines. Each node is represented by a quaternion and is connected to its three closest neighborhoods. Hence, each of the three square matrices, that compose a corner of the pattern, is indeed a 32 by 32 matrix. In the worst case, each pattern contains  $32 \cdot 32 \cdot 8 \cdot 3 = 24576$  real numbers. With a spin rate of 1 degree per minute there are 4 spin rotation a day. Over the whole mission, i.e. 2000 days, the number of non-redundant non-zero elements in the reduced normal matrix will be around

$$N \sim 8 \cdot 10^{11}. \quad (13)$$

This number of non-zero elements represents a sparsity rate of  $3 \cdot 10^{-6}$  over the whole mission considering one scalar parameters each two seconds. Let us have a look at the reduced normal equation used to test the algorithm implemented in the One Day Astronomical Solution software, ODAS [2]. This reduced normal matrix (Figure 6) was computed with one spline node per minute. The sparseness rate is, at  $8.9 \cdot 10^{-2}$ , which is disturbingly high. We are confident that the sparsity rate decreases out of the diagonal. But we would like to compute a part of the reduced normal matrix to confirm our estimate.

## 4 A direct computation

In this section, we will study the possibility to use a classical algorithm to compute the Cholesky factor of the reduced normal matrix.

### 4.1 About the minimum degree algorithm

The Cholesky factor of a positive matrix (i.e. a square matrix  $M$  such that for all vectors  $x$ , the product  $x^T M x \geq 0$ ) is a lower triangular matrix  $R$  such that  $R R^T = M$ . It is important to note that the fill-in of the Cholesky factor is not invariant to a permutation  $P$  but depends on the order of the rows and columns of the matrix  $M$ . In particular, there is a permutation  $P$  that minimizes the number of non-zero elements in the Cholesky factor. Note that, whereas the inverse of the matrix will be full, the Cholesky factor can preserve the sparseness. This is a great advantage when we want to handle large linear system. Once the Cholesky factor has been computed, the linear system  $M x = b$  can be solved considering two triangular systems.

The computation of the best permutation matrix is a NP-complete problem, i.e. a problem with a complexity that grows faster than any polynomials in size of the matrix  $M$ . It will take less resources to compute the inverse of  $M$  directly than to find the best permutation  $P$  that minimizes the Cholesky factor of  $M$ . The minimum degree algorithm is a classical heuristic process that computes a reordering of the matrix  $M$  to decrease the fill-in of the Cholesky factor. This




Figure 6: Reduced normal matrix used to test the one day astronomic solution (ODAS).

reordering is not the best, but is usually a good one. The idea of the minimum degree reordering is to consider the graph of sparseness structure and to perform at each step the Gauss pivoting on the node with the minimum number of direct connection with other nodes (see for example the description of the command *symamd* of **Matlab** in [4] and the Chapter 6.5.2 in [3]).

## 4.2 Applied to the ODAS test matrix

As we have seen in the Section 3.3 the reduced normal matrix used to test ODAS is not very sparse. Hence, the minimum degree reordering process can not be very efficient. On Figure 7, we can observe the reduced normal matrix used to test ODAS (Figure 6) reordered according to the minimum degree algorithm. On Figure 8 we observe that the Cholesky factor is almost full. Indeed the sparseness rate of the filled part of the Cholesky factor attains 0.57. Moreover, the matrix is ill-conditioned and presents a rank deficiency. As a consequence, the Cholesky computation is numerically unstable. To produce the structure of the Cholesky factor we have to add  $10^6$  on each element of the diagonal.

## 4.3 On a simplified reduced normal matrix

In order to estimate the feasibility of a direct computation with the minimum degree algorithm on a larger time scale than one day we have to find a trick. We choose to simplify the reduced normal matrix. The idea is: If we cannot find an efficient direct algorithm for the simplified matrix, a direct solution for the original problem does not seem feasible.




Figure 7: Reordering of the reduced normal matrix used to test ODAS.




Figure 8: The structure of the Cholesky factor of the reduced normal matrix used to test ODAS.




Figure 9: Condensed reduced normal matrix on 40 days.

In this simplification, we assume that:

$\mathcal{H}_1$  On each day the pattern is strictly periodic,

$\mathcal{H}_2$  The reduced normal matrix is positive definite.

The hypothesis  $\mathcal{H}_1$  will allow us to apply an efficient first reordering of the attitude unknowns on each day in order to obtain, what I call, the pseudo reduced normal matrix. The hypothesis  $\mathcal{H}_2$ , obtained by adding large positive values on the diagonal, will allow us to easily compute the Cholesky factor.

If we assume that on each day, the periodicity of the pattern is perfect, i.e. the pattern appears periodically with a period fixed by the spinning rate, we could reorder the unknown attitude parameters grouping them together periodically. For example, when we have 1000 attitude unknowns per days and if the period is 125, we consider the permutation

$$\sigma = 1, 1+125, \dots, 1+7 \cdot 125, 2, \dots, i, i+125, \dots, i+7 \cdot 125, \dots, 125+7 \cdot 125. \quad (14)$$

If the periodicity is perfect, this process is equivalent to taking the first 125 attitude unknowns and considering them as vectors of length 8.

Using this first reordering, we obtain the pseudo reduced normal matrix over 40 days. This matrix is plotted on Figure 9, considering only the first 125 attitude unknowns of each days of the reduced normal matrix. Close inspection of this figure shows again that the pattern does not appear periodically, but is slowly moving.

The Figure 10 shows the same matrix after a classical reordering process, designed to have a sparse Cholesky factor. On Figure 11 we can observe the




Figure 10: Reordered reduced normal matrix on 20 days.




Figure 11: The Cholesky factor of the reordered reduced normal matrix on 20 days.




Figure 12: Reordered reduced normal matrix on 40 days.

structure of the Cholesky factor of the reordered normal matrix over 20 days, if of course it is positive definite. It is still quite sparse.

But the problem becomes more and more complex for each days we add. On Figure 12 we can observe that the reordering process is not efficient on the reduced normal matrix for a time interval of 40 days. Indeed, the sparseness rate of the Cholesky factor has dramatically decreased as we can observe on Figure 13.

Roughly, the sparseness rate decreases by a factor 2 when we double the number of days. Hence, once Gaia will have covered the whole sky, within approximately half a year, the Cholesky factor will be almost full considering this algorithm. Moreover, when we compare the simplified matrix with a real one, we see immediately that our simplification is strong: the fill-in rate are not equal. Indeed, in our simplification we only consider the correlation between attitudes nodes due to sources. Whereas for the real matrix we have to take in account the correlation due to the splines smoothness: each observation equation (2) involves 4 splines nodes. Furthermore, as explained in the introduction, we have deliberately ignored the calibration unknowns and the “global” unknowns of the full Gaia adjustment problem. Both groups of unknowns will add more complexity and more fill-in.

## 5 Conclusion

From this data analysis, it appears that the reduced normal matrix is sparse and presents a nice structure correlated to the scanning law. But classical direct methods seem not suitable to solve the global least-square problem associated




Figure 13: Reordered reduced normal matrix on 40 days.

to the Gaia observations. In this attempt for a direct solution to the Gaia least-square problem, we consider a simpler matrix, the pseudo reduced matrix. Finding the minimum fill-in Cholesky factor of this matrix can give a lower bound on the computational complexity of the direct solution. We have seen that the minimum degree algorithm is not efficient to solve this problem. These findings make it unlikely that any direct method to solve the overall Gaia astrometric adjustment problem will be feasible. Nonetheless, this fill-up is a good property considering the geodetic problem. The solution will be accurate due to this strong correlation between the unknowns.

## References

- [1] GAIA — Composition, Formation and Evolution of the Galaxy. Technical report, July 2000. Concept and Technology Study Report, ESA-SCI(2000)4.
- [2] H.H. Bernstein, U. Bastian, and S. Hirte. First look - description of the ring solution, November 2005.
- [3] Åke Björck. *Numerical methods for least squares problems*. SIAM, 1996.
- [4] Timothy A. Davis. *Direct Methods for Sparse Linear Systems*. SIAM, 2006.
- [5] ESA. *The Hipparcos and Tycho Catalogues*. 1997. ESA SP-1200.
- [6] L. Lindegren. The scanning law for GAIA., February 1998.
- [7] L. Lindegren. A comment on the direct global solution. draft, August 2007.

- [8] L. Lindegren. Distributed processing of the attitude update through segmented spline calculations, March 2007.